Noname manuscript No.
(will be inserted by the editor)

Overlapping Cover Local Regression Machines

Mohamed Elhoseiny · Ahmed Elgammal

arXiv:1701.01218v1 [cs.LG] 5 Jan 2017

Received: date / Accepted: date

Abstract We present the Overlapping Domain Cover (ODC) notion for kernel machines, as a set of overlapping subsets of the data that covers the entire training set and optimized to be spatially cohesive as possible. We show how this notion benefit the speed of local kernel machines for regression in terms of both speed while achieving while minimizing the prediction error. We propose an efficient ODC framework, which is applicable to various regression models and in particular reduces the complexity of Twin Gaussian Processes (TGP) regression from cubic to quadratic. Our notion is also applicable to several kernel methods (*e.g*. Gaussian Process Regression(GPR) and IWTGP regression, as shown in our experiments). We also theoretically justified the idea behind our method to improve local prediction by the overlapping cover. We validated and analyzed our method on three benchmark human pose estimation datasets and interesting findings are discussed.

1 Introduction

Estimation of a continuous real-valued or a structured-output function from input features is one of the critical problems that appears in many machine learning applications. Examples include predicting the joint angles of the human body from images, head pose, object viewpoint, illumination direction, and a person's age and gender. Typically, these problems are formulated by a regression model. Recent advances in structure regression encouraged researchers to adopt it for formulating various problems with high-dimensional output spaces, such as segmentation, detection, and image reconstruction, as regression problems. However, the computational complexity of the state-of-the-art regression algorithms limits their applicability for big data. In particular, kernel-based regression algorithms such as Ridge Regression [12], Gaussian Process Regression (GPR) [18], and the Twin Gaussian Processes (TGP) [2] require inversion of kernel matrices ($O(N^3)$, where N is the number of the training points), which limits their applicability for big data. We refer to these non-scalable versions of GPR and TGP as full-GPR and full-TGP, respectively.

Khandekar et. al. [13] discussed properties and benefits of overlapping clusters for minimizing the conductance from spectral perspective. These properties of overlapping clusters also motivate studying scalable local prediction based on overlapping kernel machines. Figure 1 illustrates the notion by starting from a set of points, diving them into either disjoint and overlapping subsets, and finally learning a kernel prediction function on each (i.e., $f_i(x^*)$ for subset i, x^* is testing point). *In summary, the main question, we address in this paper, is how local kernel machines with overlapping training data could help speedup the computations and gain accurate predictions. We achieved considerable speedup and good performance on GPR, TGP, and IWTGP (Importance Weighted TGP) applied to 3D pose estimation datasets. To the best of our knowledge, our framework is the first to achieve quadratic prediction complexity for TGP. The ODC concept is also novel in the context of kernel machines and is shown here to be successfully applicable to multiple kernel-machines*. We studies in this work GPR and TGP and IWTGP (a third model) kernel machines. The remainder of this paper is organized as follows: Section 2 and 4 presents some motivating kernel machines and the related work. Section 5 presents our approach and a theoretical justification for our ODC concept. Section 6 and 7 presents our experimental validation and conclusion.

Mohamed Elhoseiny[1] and Ahmed Elgammal[2]
[1]Facebook AI Research
[2]Department of Computer Science, Rutgers University
E-mail: elhoseiny@fb.com, elgammal@cs.rutgers.edu

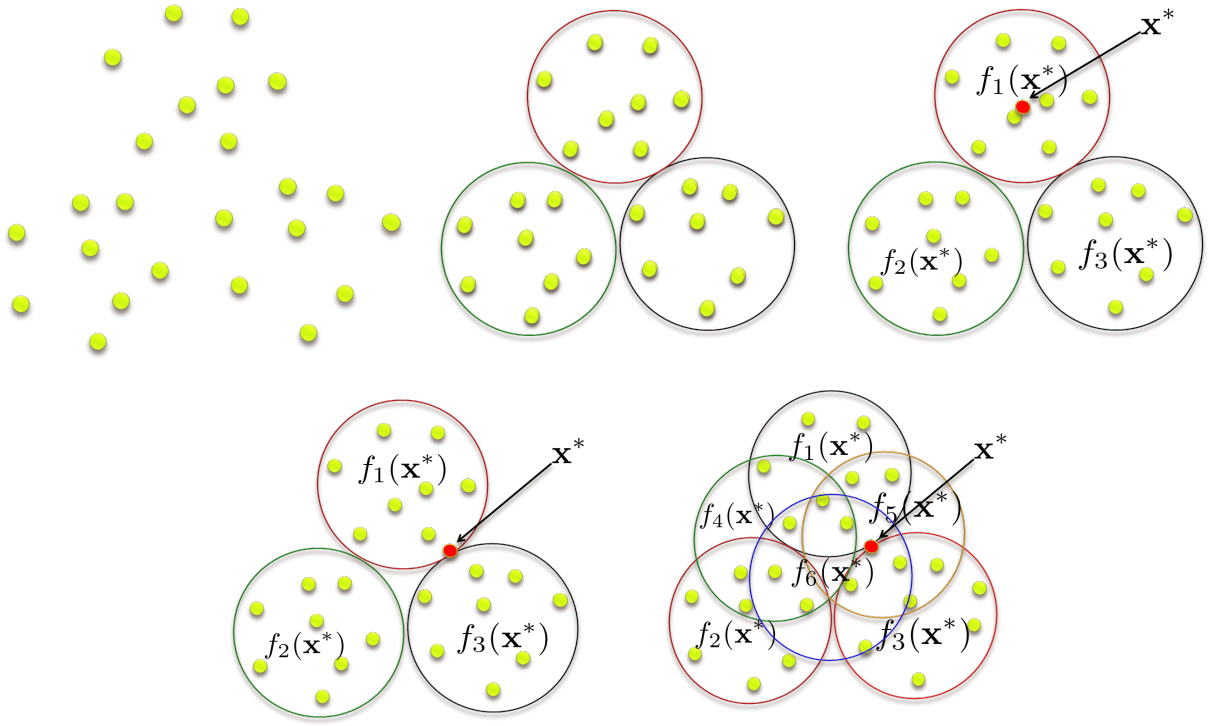

Fig. 1: Top: Left:24 points, Middle: Overlapping Cover, Right: disjoint kernel machines of 8 points (evaluating x^* near a middle of a kernel machine). **Bottom**: Left: disjoint kernel machine evaluation on boundary), Right: 6 Overlapping kernel machines of 8 points. $f_i(\mathbf{x}^*)$ is the i^{th} kernel machine prediction for $\mathbf{x}^*$ test point.

2 Background on Full GPR and TGP Models

In this section, we show example kernel machines that motivated us to propose the ODC framework to improve their performance and scalability. Specifically, we review GPR for single output regression, and TGP for structured output regression. We selected GPR and TGP kernel machines for their increasing interest and impact. However, our framework is not restricted to them.

GPR [18] assumes a linear model in the kernel space with Gaussian noise in a single-valued output, i.e., $y = f(\mathbf{x}) + \mathcal{N}(0, \sigma_n^2)$, where $\mathbf{x} \in \mathbb{R}^{d_X}$ and $y \in \mathbb{R}$. Given a training set $\{\mathbf{x}_i, y_i, i = 1 : N\}$, the posterior distribution of y given a test point $\mathbf{x}_*$ is:

$$\begin{aligned} p(y|\mathbf{x}_*) = \mathcal{N}(\boldsymbol{\mu}_y = \mathbf{k}(\mathbf{x}_*)^\top (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{f}, \\ \sigma_y^2 = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{x}_*)^\top (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}(\mathbf{x}_*)) \end{aligned} \tag{1}$$

where $k(\mathbf{x}, \mathbf{x}')$ is kernel defined in the input space, $\mathbf{K}$ is an $N \times N$ matrix, such that $\mathbf{K}(l, m) = k(\mathbf{x}_l, \mathbf{x}_m)$, $\mathbf{k}(\mathbf{x}_*) = [k(\mathbf{x}_*, \mathbf{x}_1), , ..., k(\mathbf{x}_*, \mathbf{x}_N)]^\top$, $\mathbf{I}$ is an identity matrix of size N, σ_n is the variance of the measurement noise, $\mathbf{f} = [y_1, \cdots, y_N]^\top$. GPR could predict structured output $\mathbf{y} \in \mathbb{R}^{d_Y}$ by training a GPR model for each dimension. However, this indicates that GPR does not capture dependency between output dimensions which limit its performance.

TGP [2] encodes the relation between both inputs and outputs using GP priors. This was achieved by minimizing the Kullback-Leibler divergence between the marginal GP of outputs (e.g., poses) and observations (e.g., features). Hence, TGP prediction is given by:

$$\begin{aligned} \hat{\mathbf{y}}(\mathbf{x}_*) = & \underset{\mathbf{y}}{\operatorname{argmin}} [k_Y(\mathbf{y}, \mathbf{y}) - 2\mathbf{k}_Y(\mathbf{y})^\top (\mathbf{K}_X + \lambda_X \mathbf{I})^{-1} \mathbf{k}_X(\mathbf{x}_*) \\ & - \eta log(k_Y(\mathbf{y}, \mathbf{y}) - \mathbf{k}_Y(y)^\top (\mathbf{K}_Y + \lambda_Y \mathbf{I})^{-1} \mathbf{k}_Y(\mathbf{y}))] \end{aligned} \tag{2}$$

where $\eta = k_X(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_X(\mathbf{x}_*)^\top (\mathbf{K}_X + \lambda_X \mathbf{I})^{-1} \mathbf{k}_X(\mathbf{x}_*)$, $k_X(\mathbf{x}, \mathbf{x}') = exp(\frac{-\|\mathbf{x}-\mathbf{x}'\|}{2\rho_x^2})$ and $k_Y(\mathbf{y}, \mathbf{y}') = exp(\frac{-\|\mathbf{y}-\mathbf{y}'\|}{2\rho_y^2})$ are Gaussian kernel functions for input feature $\mathbf{x}$ and output vector $\mathbf{y}$, ρ_x and ρ_y are the kernel bandwidths for the input and the output . $\mathbf{k}_Y(\mathbf{y}) = [k_Y(\mathbf{y}, \mathbf{y}_1), ..., k_Y(\mathbf{y}, \mathbf{y}_N)]^\top$, where N is the number of the training examples. $\mathbf{k}_X(\mathbf{x}_*) = [k_X(\mathbf{x}_*, \mathbf{x}_1), ..., k_X(\mathbf{x}_*, \mathbf{x}_N)]^\top$, and λ_X and λ_Y are regularization parameters to avoid overfitting. This optimization problem can be solved using a quasi-Newton optimizer with cubic polynomial line search [2]; we denote the number of steps to convergence as l_2.

Table 1: Comparison of computational Complexity of training and testing for each of Full, NN (Nearest Neighbor), FITC, Local-RPC, and our ODC. Training is the time include all computations that does not depend on test data, which includes clustering in some of these methods. Testing includes computations only needed for prediction

	Training for GPR and TGP			Testing for each point		
	Ekmeans Clustering	**RPC Clustering**	**Model training**	**GPR-Y**	**GPR-Var**	**TGP-Y**
Full	-	-	$O(N^3 + N^2 d_X)$	$O(N \cdot (d_X + d_Y)$	$O(N^2 \cdot d_Y)$	$O(l_2 \cdot N^2 \cdot d_Y)$
NN [2]	-	-	-	$O(M^3 \cdot d_Y)$	$O(M^3 \cdot d_Y)$	$O(M^3 + l_2 \cdot M^2 \cdot d_Y)$
FIC (GPR only, $d_Y = 1$ [22])	-	-	$O(M^2 \cdot (N + d_X))$	$O(M \cdot d_X)$	$O(M^2)$	-
Local-RPC (only GPR, $d_Y = 1$ [5])	-	$N \cdot log(\frac{N}{M})$	$O(M^2 \cdot (N + d_X))$	$O(M \cdot d_X)$	$O(M^2)$	-
ODC (our framework)	$O(N \cdot \frac{N}{(1-p)M} \cdot d_X \cdot l_1)$	$O(N \cdot log(\frac{N}{(1-p)M}) \cdot d_X)$	$O(M^2 \cdot (\frac{N}{1-p} + d_X))$	$O(K' \cdot M \cdot (d_X + d_Y))$	$O(K' \cdot M^2 \cdot d_Y)$	$O(l_2 \cdot K' \cdot M^2 \cdot d_Y)$

3 Importance Weighted Twin Gaussian Processes (IWTGP)

Yamada et al [26] proposed the importance-weighted variant of twin Gaussian processes [2] called IWTGP. The weights are calculated using RuLSIF [27] (relative unconstrained least-squares importance fitting). The weights were modeled as $w_\alpha(\mathbf{x}, \boldsymbol{\theta}) = \sum_{l=1}^{n_{te}} \theta_l k(\mathbf{x}, \mathbf{x}_l)$ to minimize $E_{p_{te}(x)}[\,(w_\alpha(x,\theta) - w_\alpha(\mathbf{x}))^2\,]$. where $k(\mathbf{x}, \mathbf{x}_l) = exp(-\frac{\|\mathbf{x}-\mathbf{x}_l\|}{2\tau^2})$, $w_\alpha(\mathbf{x}) = \frac{p_{te}(\mathbf{x})}{(1-\alpha)p_{te}(\mathbf{x}) + \alpha p_{tr}(\mathbf{x})}$, $0 \leq \alpha \leq 1$. To cope with this instability issue, setting α to $0 \leq \alpha \leq 1$ is practically useful for stabilizing the covariate shift adaptation, even though it cannot give an unbiased model under covariate shift [27]. According [26] the optimal $\hat{\boldsymbol{\theta}}$ vector is computed in a closed form solution as follows.to

$$\hat{\boldsymbol{\theta}} = (\hat{\mathbf{H}} + \nu \mathbf{I})^{-1} \hat{\mathbf{h}} \tag{3}$$

where $\hat{\mathbf{H}}_{l,l'} = \frac{1-\alpha}{n_{te}} \sum_{i=1}^{n_{te}} k(\mathbf{x}_i^{te}, \mathbf{x}_l^{te}) k(\mathbf{x}_i^{te}, \mathbf{x}_{l'}^{te}) + \frac{\alpha}{n_{tr}} \sum_{j=1}^{n_{tr}} k(\mathbf{x}_j^{tr}, \mathbf{x}_l^{te}) k(\mathbf{x}_j^{tr}, \mathbf{x}_{l'}^{te})$, $\hat{\mathbf{h}}$ is an n_{te}- dimensional vector with the l^{th} element $\hat{\mathbf{h}}_l = \frac{1}{n_{te}} \sum_{i=1}^{n_{te}} k(\mathbf{x}_i^{te}, \mathbf{x}_l^{te})$, $\mathbf{I}$ is an $n_{te} \times n_{te}$-dimensional identity matrix. where n_{te} and n_{tr} and the number of testing and training points respectively. Model selection of RuLSIF is based on cross-validation with respect to the squared-error criterion J in [27]. Having computed $\hat{\boldsymbol{\theta}}$, each input and output examples are simply re-weighted by $w_\alpha^{\frac{1}{2}}$ [26]. Therefore, the output of the importance weighted TGP (IWTGP) is given by

$$\hat{y} = \underset{y}{\operatorname{argmin}} [K_Y(\mathbf{y}, \mathbf{y}) - 2k_y(\mathbf{y})^T \mathbf{u}_w - \eta_w log(K_Y(\mathbf{y}, \mathbf{y}) - k_y(\mathbf{y})^T \mathbf{W}^{\frac{1}{2}} (\mathbf{W}^{\frac{1}{2}} \mathbf{K}_Y \mathbf{W}^{\frac{1}{2}} + \lambda_y I)^{-1} \mathbf{W}^{\frac{1}{2}} k_y(\mathbf{y}))] \tag{4}$$

where $\mathbf{u}_w = \mathbf{W}^{\frac{1}{2}} (\mathbf{W}^{\frac{1}{2}} \mathbf{K}_X \mathbf{W}^{\frac{1}{2}} + \lambda_x I)^{-1} \mathbf{W}^{\frac{1}{2}} k_x(\mathbf{x})$, $\eta_w = k_X(\mathbf{x}, \mathbf{x}) - k_x(\mathbf{x})^T \mathbf{u}_w$. Similar to TGP, IWTGP can also be solved using a second order, BFGS quasi-Newton optimizer with cubic polynomial line search for optimal step size selection.

Table 1 shows the training an testing complexity of full GPR and TGP models, where d_Y is the dimensionality of the output. Table 1 also summarizes the computational complexity of the related approximation methods, discussed in the following section, and our method. N.

4 Related Work on Approximation Methods

Various approximation approaches have been presented to reduce the computational complexity in the context of GPR. As detailed in [16], approximation methods on Gaussian Processes may be categorized into three trends: matrix approximation, likelihood approximation, and localized regression. The matrix approximation trend is inspired by the observation that the kernel matrix inversion is the major part of the expensive computation, and thus, approximating the matrix by a lower rank version, $M \ll N$ (e.g., Nyström Method [25]). While this approach reduces the computational complexity from $O(N^3)$ to $O(NM^2)$ for training, there is no guarantee on the non-negativity of the predictive variance [18]. In the second trend, likelihood approximation is performed on testing and training examples, given M artificial examples known as inducing inputs, selected from the training set (e.g., Deterministic Training Conditional (DTC) [19], Full Independent conditional (FIC) [22], Partial Independent Conditional (PIC) [21]). The drawback of this trend is the dilemma of selecting M inducing points, which might be distant from the test point, resulting in a performance decay; see Table 1 for the complexity of FIC.

A third trend, localized regression, is based on the belief that distant observations are almost unrelated. The prediction of a test point is achieved through its M nearest points. One technique to implement this notion is through decomposing the training points into disjoint clusters during training, where prediction functions are learned for each of them [16]. At test time, the prediction function of the closest cluster is used to predict the corresponding output. While this method is efficient, it introduces discontinuity problems on boundaries of the subdomains. Another way to implement local regression is through Mixture of Experts (MoE) as an Ensemble method to make prediction based on computing the final output by combining outputs of local predictors called experts (see a study on MoE methods [28]). Examples include Bayesian committee machine (BCM [23]), local probabilistic regression (LPR [24]), mixture of Tree of Gaussian Processes (GPs) [9], and Mixture of GPs [18]. While these approaches overcome the discontinuity problem by the combination mechanism, they suffer from intensive complexity at test time, which limits its applicability in large-scale setting, e.g., Tree of GPs and Mixture of GPs,

involve complicated integration, approximated by computationally expensive sampling or Monte Carlo simulation.

Park etal. [16] proposed a large-scale approach for GPR by domain decomposition on up to 2D grid on input, where a local regression function is inferred for each subdomain such that they are consistent on boundaries. This approach obviously lacks a solution to high-dimensional input data because the size of the grid increases exponentially with the dimensions, which limits its applicability. More recently, [5] proposed a Recursive Partitioning Scheme (RPC) to decompose the data into non-overlapping equal-size clusters, and they built a GPR on each cluster. They showed that this local scheme gives better performance than FIC [22] and other methods. However, this partitioning scheme obviously lacks consistency on the boundaries of the partitions and it was restricted to single-output GPR. Table 1 shows the complexity of this scheme denoted by local-RPC for GPR.

Beyond GPR, we found that local regression was adopted differently in structured regression models like Twin Gaussian Processes (TGP) [2], and also an data bias version of it, denoted by IWTGP [26]. TGP and IWTGP outperform not only GPR in this task, but also various regression models including Hilbert Schmidt Independence Criterion (HSIC) [10], Kernel Target Alignment (KTA) [6], and Weighted-KNN [18]. Both TGP and IWTGP have no closed-form expression for prediction. Hence, the prediction is made by gradient descent on a function that needs to compute the inverse of both the input and output kernel matrices, $O(N^3)$ complexity. Practically, both approaches have been applied by finding the $M \ll N$ Nearest-Neighbors (NN) of each test point in [2] and [26]. The prediction of a test point is $O(M^3)$ due to the inversion of $M \times M$ input and output kernel Matrices. However, NN scheme has three drawbacks: (1) A regression model is computed for each test point, which results in a scalability problems in prediction (i.e., Matrix inversions on the NN of each each test point), (2) Number of neighbors might not be large enough to create an accurate prediction model since it is constrained by the first drawback, (3) It is inefficient compared with the other schemes used for GPR. Table 1 shows the complexity of this NN scheme.

5 ODC Framework

The problems of the existing approaches, presented above, motivated us to develop an approach that satisfies the properties listed in table 2. The table also shows which of these properties are satisfied for the relevant methods. In order to satisfy all the properties, we present the Overlapping Domain Cover (ODC) notion. We define the ODC as a collection of overlapping subsets of the training points, denoted by subdomains, such that they are as spatially coherent as possible. During training, an ODC is computed such that each subdomain overlaps with the neighboring subdomains.

Table 2: Contrast against most relevant methods

	[16]	FIC/PIC [22]	NN [2]	ODC
Accurate	No for high input dimension	No	Yes	Yes
Efficient	No	Yes	No	Yes
Scalable to arbitrary input dimension	No (2D)	Yes	Yes	Yes
Consistent on Boundaries	Yes	No	Yes	Yes
supported kernel machines	GPR	GPR	TGP	GPR, TGP, IWTGP and others
Easy to parallelize	No	No	Yes	Yes

Then, a local prediction model (kernel machine) is created for each subdomain and the computations that does not depend on the test data are factored out and precomputed (e.g. inversion of matrices). The nature of the ODC generation makes these kernel machines consistent in the overlapped regions, which are the boundaries since we constraint the subdomains to be coherent. This is motivated by the notion that data lives on a manifold with local properties and consistent connections between its neighboring regions. On prediction, the output is calculated as a reduction function of the predictions on the closed subdomain(s). Table 1 (the last row) shows the complexity for our generalized ODC framework, detailed in Sec 5.1 and 5.2. In contrast to the prior work, our ODC framework is designed to cover structured regression setting, $d_Y > 1$ and to be applicable to GPR, TGP, and many other models.

Notations. Given a set of input data $X = \{\mathbf{x}_1, \cdots, \mathbf{x}_N\}$, our prediction framework firstly generates a set of non-overlapping equal-size partitions, $C = \{C_1, \cdots, C_K\}$, such that $\cup_i C_i = X$, $|C_i| = N/K$. Then, the ODC is defined based on them as $\mathcal{D} = \{D_1, \cdots, D_K\}$, such that $|D_i| = M \, \forall i$, $D_i = C_i \cup O_i, \forall i$. O_i a the set of points that overlaps with the other partitions, i.e., $O_i = \{x : x \in \{\cup_{j \neq i} C_j\}\}$, such that $|O_i| = p \cdot M$, $|C_i| = (1-p) \cdot M$, $0 \leq p \leq 1$ is the ratio of points in each overlapping subdomain, D_i, that belongs to/overlaps with partitions, other than its own, C_i.

It is important to note that, the ODC could be specified by two parameters, M and p, which are the number of points in each subdomain and the ratio of overlap respectively; this is since $K = N/(1-p)M$. This parameterization of ODC generation is reasonable for the following reasons. First, M defines the number of points that are used to train each local kernel machine, which controls the performance of the local prediction. Second, given M and that $K = N/(1-p)M$, p defines how coarse/fine the distribution of kernel machines are. It is not hard to see that as p goes to 0, the generated ODC reduces to the set of non-overlapping clusters. Similarly, as p approaches $1 - 1/M$, the ODC reduces to generating a cluster at each point with maximum overlap with other clusters, i.e., $K = N$, $|C_i| = 1$, and $|O_i| = M - 1$. *Our main claim is two fold. First, precomputing local kernel machines (e.g. GPR, TGP, IWTGP) during training on the ODC significantly increase the speedup on prediction time. Second, given a fixed M and N, as p increases, local prediction per-*

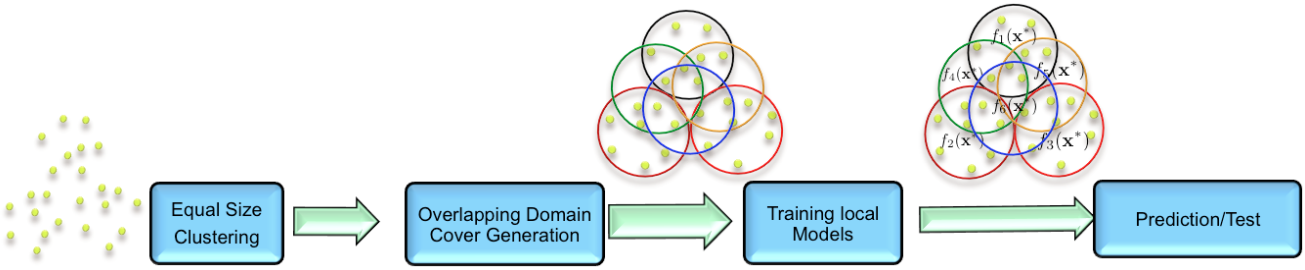

Fig. 2: ODC Framework

formance increases, theoretically supported by Lemma 51

Lemma 51. Under ODC notion, as the overlap p increases, the closer the nearest model to an arbitrary test point and the more likely that model get trained on a big neighborhood of the test point.

Proof. We start by outlining the main idea behind the proof, which is directly connected to the fact that $K = N/(1-p)M$, which indicates that the number of local models increases as p increases given fixed N and M. Under the assumption that the local models are spatially cohesive, $p \to 1$ theoretically indicates that there is a local model centered at each point in the space (*i.e.* $K = \infty$). Hence, as p increases, the distribution of the kernel machines is the finest and the more likely a test point to find the closest kernel machines trained on a big neighborhood of it leading to more accurate prediction. Meanwhile, as p goes to 0, the distribution is the coarsest and the less likely a test point finds, the closest kernel machines, trained on a big neighborhood.

Let's assume that each kernel machine is defined on M points that are spatially cohesive, covering the space of N points with $\frac{N}{(1-p)M}$. Let's assume that center of the M points in kernel machine i is μ_i, the the Co-variance matrix of these points are Σ_i. Hence

$$\begin{aligned} p(\mathbf{x}|D_i) &= \mathcal{N}(\mu_i, \Sigma_i) \\ &= (2\pi)^{-\frac{d_X}{2}} |\Sigma_i|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{x}-\mu_i)^\intercal \Sigma_i^{-1}(\mathbf{x}-\mu_i)} \end{aligned} \tag{5}$$

where $\mathcal{N}(\mu_i, \Sigma_i)$ is a normal distribution of mean μ_i and Co-variance matrix Σ_i.

Let's assume that there are two ODCs, ODC_1 and ODC_2, defined on the same N points, the first one has overlap p_1 and the second one is with overlap p_2, such that, $p_2 > p_1$. Let's assume that the number of kernel machines in ODC_1 and ODC_2 are K_1 and K_2, respectively. Hence,

$$K_1 = \frac{N}{(1-p_1)M} \quad , \quad K_2 = \frac{N}{(1-p_2)M} \tag{6}$$

Since $p_2 > p_1$, $0 \leq p_1 < 1$ and $0 \leq p_2 < 1$, then $K_2 > K_1$, which indicates that the number of kernel machines in ODC_2 with higher overlap is bigger than the number of kernel machines in ODC_2. Let's assume that there is an test point $\mathbf{x}^*$ and define that the probability that $\mathbf{x}^*$ is captured by the ODC to be proportional to the maximum probability of $\mathbf{x}^*$ among the domains.

$$\begin{aligned} p(\mathbf{x}^*) &= \sum_{i=1}^{K} p(\mathbf{x}^*, D_i) \\ &= \sum_{i=1}^{K} p(\mathbf{x}^*|D_i)\delta(p(\mathbf{x}^*|D_i) - max_{j=1}^{K}(p(\mathbf{x}^*|D_i))) \\ &= max_{i=1}^{K} p(\mathbf{x}^*|D_i) \\ &= (2\pi)^{-\frac{d_X}{2}} max_{i=1}^{K} |\Sigma_i|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{x}^*-\mu_i)^\intercal \Sigma_i^{-1}(\mathbf{x}^*-\mu_i)} \end{aligned} \tag{7}$$

where $\delta(0) = 1$, 0 otherwise. The reason behind this definition of $p(x^*)$ is that our method select the domain of preduction based on $argmax_{i=1}^{K} p(\mathbf{x}^*|D_i)$. Hence $p_{ODC_1}(x^*) = max_{i=1}^{K_1} p_{ODC_1}(\mathbf{x}^*|D_i)$ and $p_{ODC_2}(x^*) = max_{i=1}^{K_2} p_{ODC_2}(\mathbf{x}^*|D_i)$.

We start by the case where the points are uniformally distributed in the space. Under this condition and assuming that spatially cohesive domain cover, this leads to that $p(\mathbf{x}^*|D_i) \approx \mathcal{N}(\mu_i, \Sigma) \forall i$, where $\Sigma_1 = \Sigma_2 \cdots = \Sigma_K = \Sigma$. Hence

$$\begin{aligned} p(\mathbf{x}^*|D_i) &\propto e^{-\frac{1}{2}(\mathbf{x}^*-\mu_i)^\intercal \Sigma^{-1}(\mathbf{x}^*-\mu_i)} \\ ln(p(\mathbf{x}^*|D_i)) &\propto -(\mathbf{x}^*-\mu_i)^\intercal \Sigma^{-1}(\mathbf{x}^*-\mu_i) \end{aligned} \tag{8}$$

Then

$$\begin{aligned} p(\mathbf{x}^*) &= max_{i=1}^{K} p(\mathbf{x}^*|D_i) \\ &= (2\pi)^{-\frac{d_X}{2}} \Sigma|^{-\frac{1}{2}} max_{i=1}^{K} |e^{-\frac{1}{2}(\mathbf{x}-\mu_i)^\intercal \Sigma^{-1}(\mathbf{x}-\mu_i)} \\ &\propto max_{i=1}^{K} e^{-\frac{1}{2}(\mathbf{x}-\mu_i)^\intercal \Sigma^{-1}(\mathbf{x}-\mu_i)} \\ ln(p(\mathbf{x}^*)) &\propto max_{i=1}^{K} - (\mathbf{x}-\mu_i)^\intercal \Sigma^{-1}(\mathbf{x}-\mu_i) \end{aligned} \tag{9}$$

Hence, $p(\mathbf{x}^*)$ gets maximized as it get closer to one of the centers of the domains μ_i, defined by the ODC. It is not hard to seen that that chances of $\mathbf{x}^*$ to be closer to one of the centers covered by ODC_2 is higher than ODC_2, especially when $p_2 \gg p_1$. This is since $K_1 = \frac{N}{(1-p_1)M}, K_2 =$

$\frac{N}{(1-p_2)M}$. Hence $K_2 \gg K_1$ when $p_2 \gg p_1$. For instance, when $p_1 = 0$ and $p2 = 0.9$, this leads to that ODC_1 will generate $K_1 = \frac{N}{M}$ domains, while ODC_2 will generate $K_2 = \frac{10 \cdot N}{M} = 10K_1$, which is ten times more domains and centers. The fact that there are much more domains if $K_2 \gg K_1$ together with that there domains are spatially cohesive leads to $max_{i=1}^{K_1} - (\mathbf{x}^* - \mu_i^1)^\mathsf{T} \Sigma_1^{-1} (\mathbf{x}^* - \mu_i^1) \gg max_{i=1}^{K_2} - (\mathbf{x}^* - \mu_i^2)^\mathsf{T} \Sigma_2^{-1} (\mathbf{x}^* - \mu_i^2)$. The proof of this statement derives from the fact that $max_{i=1}^{K} - (\mathbf{x}^* - \mu_i)^\mathsf{T} \Sigma^{-1} (\mathbf{x}^* - \mu_i)$ is could maximized by (1) if $\mathbf{x}^*$ gets very close to one of $\mu_i, i = 1 : K$,and (2) smaller variance $|\Sigma|$, which is minimized by the nature by which ODC is created, since each domain i is created by neighboring points to its center (*i.e.* $|\Sigma_1| \gg |\Sigma_2|$). This directly leads to that if $K_2 \gg K_1$ then $max_{i=1}^{K_1} - (\mathbf{x}^* - \mu_i^1)^\mathsf{T} \Sigma_1^{-1} (\mathbf{x}^* - \mu_i^1) \gg max_{i=1}^{K_2} - (\mathbf{x}^* - \mu_i^2)^\mathsf{T} \Sigma_2^{-1} (\mathbf{x}^* - \mu_i^2)$. Hence, $p_{ODC_2}(x^*) \gg p_{ODC_1}(x^*)$.

Even if the points are not uniformally distributed, it is still more likely that an ODC with higher overlap would have higher $p(x^*)$, since x^* is close under expectation to one of the centers if more spatially cohesive domains are generated which increases with higher overlap. Our experiments also proves that the ODC concept generalizes on three real dataset where the training points are not distributed uniformally. □

5.1 Training

There are several overlapping clustering methods that include (e.g. [17] and [3]), which looks relevant for our framework. However these methods does not fit our purpose both equal-size constraints for the local kernel machines. We also found them very slow in practice because their complexity varies from cubic to quadratic (with a big constant factor) on the training-set. These problems motivated us to propose a practical method that builds overlapping local kernel-machines with spatial and equal-size constraints. These constraints are critical for our purpose since the number of points in each kernel-machine determine its local performance. Hence, our training phase is two steps: (1) the training data is split into $K = N/(1-p)M$ equal-sized clusters of $(1-p)M$ points. (2) an ODC with K overlapping subdomains is generated by augmenting each cluster with $p \cdot M$ points from the neighboring clusters.

5.1.1 Equal-size Clustering

There are recent algorithms that deal with size constraints in clustering. For example, [29] formulated the problem of clustering with size constraints as a linear programming problem. However such algorithms are not computationally efficient, especially for large scale datasets (e.g., Human3.6M). We study two efficient ways to generate equal size clusters; see Table 1 (last row) for their ODC-complexity.

Recursive Projection Clustering (RPC) [5]. In this method, the training data is partitioned to perform GPR prediction. Initially all data points are put in one cluster. Then, two points are chosen randomly and orthogonal projection of all the data onto the line connecting them is computed. Depending on the median value of the projections, The data is then split into two equal size subsets. The same process is then applied to each cluster to generate 2^l clusters after l repetitions. The iterations stops once $2^l > K$. As indicated, the number of clusters in this method has to be a power of two and it might produce long thin clusters.

Equal-Size K-means (EKmeans). We propose a variant of k-means clustering [11] to generate equal-size clusters. The goal is to obtain disjoint partitioning of X into clusters $C = \{C_1, \cdots, C_K\}$, similar to the k-means objective, minimizing the within-cluster sum of squared Euclidean distances, $C = \arg_C J(C) = \min \sum_{j=1}^{K} \sum_{\mathbf{x_i} \in C_j} d(\mathbf{x_i}, \boldsymbol{\mu}_j)$, where $\boldsymbol{\mu}_i$ is the mean of cluster C_i, and $d(\cdot, \cdot)$ is the squared distance. Optimizing this objective is NP-hard and k-means iterates between the assignment and update steps as a heuristic to achieve a solution; l_1 denotes number of iterations of kmeans. We add equal-size constraints $\forall (1 \leq i \leq K), |C_i| = N/K = (1-p)M$.

In order to achieve this partitioning, we propose an efficient heuristic algorithm, denoted by *Assign and Balance (AB) EKmeans*. It mainly modifies the assignment step of the k-means to bound the size of the resulting clusters. We first assign the points to their closest see center as typically done in the assignment step of k-means. We use $C(\mathbf{x}_p)$ to denote the cluster assignment of a given point $\mathbf{x}_p$. This results in three types of clusters: balanced, overfull, and underfull clusters. Then some of the points in the overfull clusters are redistributed to the underfull clusters by assigning each of these points to the closest underfull cluster. This is achieved by initializing a pool of overfull points defined as $\tilde{X} = \{\mathbf{x}_p : \mathbf{x}_p \in C_l, |C_l| > N/K\}$; see Figure 3.

Let us denote the set of underfull clusters by $\tilde{C} = \{C_p : |C_p| < N/K\}$. We compute the distances $d(\mathbf{x}_i, \boldsymbol{\mu}_j), \forall \mathbf{x_i} \in \tilde{X}$ and $C_i \in \tilde{C}$. Iteratively, we pick the minimum distance pair $(\mathbf{x}_p, \boldsymbol{\mu}_l)$ and assign $\mathbf{x}_p$ to cluster C_l instead of cluster $C(\mathbf{x}_p)$. The point is then removed from the overfull pool. Once an underfull cluster becomes full it is removed from the underfull pool, once an overfull cluster is balanced, the remaining points of that cluster are removed from overfull pool. The intuition behind this algorithms is that, the cost associated with the initial optimal assignment (given the computed means) is minimally increased by each swap since we pick the minimum distance pair in each iteration. Hence the cost is kept as low as possible while balancing the clusters. We denote the the name of this Algoirthm as Assign and Balance EKmeans. Algorithm 1 illustrates the overall assignment step and Fig. 4 visualizes the balancing step.

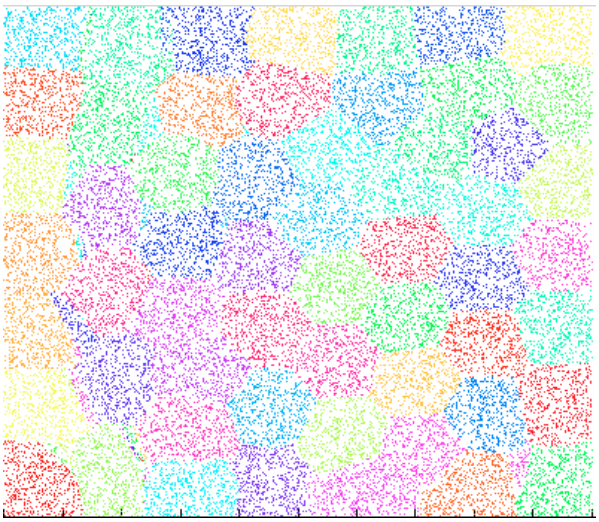

Fig. 3: AB-EKmeans on 300,000 2D points, K= 57

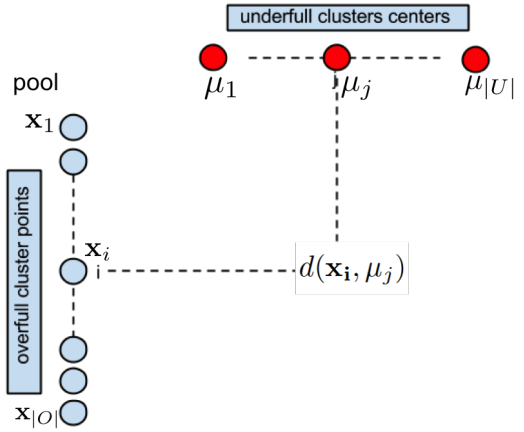

Fig. 4: AB Kmeans: Balancing Step

Input: $\mathbf{X}(N \times d_x), \{\boldsymbol{\mu}_i\}_{i=1}^{K}$
Output: labels
1- Assign the points initially to its closest center; this will put the clusters into 3 groups (1) balanced clusters (2) overflowed clusters (3) under-flowed clusters.
2- Create a matrix $\mathbf{D} \in R^{N \times K}$, where $\mathbf{D}[i, j]$ is the distance between the i^{th} point to the j^{th} cluster center; rows are restricted points belongs only to the overflowed clusters; columns are restricted to underflowed cluster centers
3- Get the coordinate (i_*, j_*) that maps the smallest distance in $\mathbf{D}$.
4- Remove the i_*^{th} row from matrix $\mathbf{D}$ and mark it as assigned to the j^{th} cluster
5- If the size of the cluster j achieves the ideal size (i.e. n/K), then remove the j^{th} column from matrix $\mathbf{D}$.
6- Go to step 3 if there is still unassigned points

Algorithm 1: Assign and Balance (AB) k-means: Assignment Step

5.1.2 Overlapping Domain Cover(ODC) Model

Having generated the disjoint equal size clusters, we generate the ODC subdomains based on the overlapping ratio p, such that $p \cdot M$ points are selected from the neighboring clusters. Let's assume that we select only the closest r clusters to each cluster, C_i is closer to C_j than C_k if $\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\| < \|\boldsymbol{\mu}_i - \boldsymbol{\mu}_k\|$. It is important to note that r must be greater than $p/(1-p)$ in order to supply the required $p \cdot M$ points; this is since number of points in each cluster is $(1-p)M$. Hence, the minimum value for r is $\lceil (p \cdot M)/((1-p) \cdot M) \rceil = \lceil p/(1-p) \rceil$ clusters. Hence, we parametrize r as $r = \lceil t \cdot p/(1-p) \rceil, t \geq 1$. We study the effect of t in the experimental results section. Having computed r from p and t, each subdomain D_i is then created by merging the points in the cluster C_i with $p \cdot M$ points, retrieved from the r neighboring clusters. Specifically, the points are selected by sorting the points in each of r clusters by the distance to $\boldsymbol{\mu}_i$. The number of points retrieved for each of the r neighboring clusters is inversely proportional to the distance of its center to $\boldsymbol{\mu}_i$. If a subset of the r clusters are requested to retrieve more than its capacity (i.e., $(1-p)M$), the set of the extra points are requested from the remaining clusters giving priority to the closer clusters (i.e., starting from the nearest neighboring cluster to the cluster on which the subdomain is created). As $t = 1$ and p increases, all points that belong to the r clusters tends to be merged with C_i. In our framework, we used FLANN [15] for fast NN-retrieval; see pseudo-code of ODC generation in Appendix C.

After the ODC is generated, we compute the the sample normal distribution using the points that belong to each subdomain. Then, a local kernel machine is trained for each of the overlapping subdomains. We denote the point set normal distribution of the subdomains as $p(\mathbf{x}|D_i) = \mathcal{N}(\boldsymbol{\mu}_i' \in R^{d_X}, \Sigma_i' \in R^{d_X \times d_X})$; ${\Sigma'}_i^{-1}$ is precomputed during the training for later use during the prediction. Finally, we factor out all the computations that does not depend on the test point (for GPR, TGP, IWTGP) and store them with each sub domain as its local kernel machine. We denote the training model for subdomain i as $\mathcal{M}^i$, which is computed as follows for GPR and TGP respectively.

GPR. Firstly, we precompute $(\mathbf{K}_j^i + \sigma^2_{n_j^i}\mathbf{I})^{-1}$, where $\mathbf{K}_j^i$ is an $M \times M$ kernel matrix, defined on the input points in D_i. Each dimension j in the output could have its own hyper-parameters, which results in a different kernel matrix for each dimension $\mathbf{K}_j^i$. We also precompute $(\mathbf{K}_j^i + \sigma^2_{n_j^i}\mathbf{I})^{-1}\mathbf{y}_j$ for each dimension. Hence $\mathcal{M}^i_{GPR} = \{(\mathbf{K}_j^i + \sigma^2_{n_j^i}\mathbf{I})^{-1}, (\mathbf{K}_j^i + \sigma^2_{n_j^i}\mathbf{I})^{-1}\mathbf{y}_j), j = 1 : d_Y\}$.

TGP. The local kernel machine for each subdomain in TGP case is defined as $\mathcal{M}^i_{TGP} = \{(\mathbf{K}^i_X + \lambda^i_X\mathbf{I})^{-1}, (\mathbf{K}^i_Y + \lambda^i_Y\mathbf{I})^{-1}\}$, where $\mathbf{K}^i_X$ and $\mathbf{K}^i_Y$ are $M \times M$ kernel matrices de-

fined on the input points and the corresponding output points respectively, which belong to domain i.

IWTGP. It is not obvious how to factor out computations that does not depend on the test data in the case of IWTGP, since the computational extensive factor(i.e., $(\mathbf{W}^{i\frac{1}{2}}\mathbf{K}_X^i\mathbf{W}^{i\frac{1}{2}}+\lambda_x^i\mathbf{I})^{-1}$, $(\mathbf{W}^{i\frac{1}{2}}\mathbf{K}_Y^i\mathbf{W}^{i\frac{1}{2}}+\lambda_y^i\mathbf{I})^{-1}$) does depend on the test set since $\mathbf{W}^i$ is computed on test time. To help factor out the computation, we used linear algebra to show that

$$(\mathbf{D}\ \mathbf{A}\ \mathbf{D}+\lambda\mathbf{I})^{-1}=\mathbf{D}^{-1}\mathbf{A}^{-1}\mathbf{D}^{-1}-\frac{\lambda\mathbf{D}^{-2}\mathbf{A}^{-2}\mathbf{D}^{-2}}{1+\lambda\cdot tr(\mathbf{D}^{-1}\mathbf{A}^{-1}\mathbf{D}^{-1})} \tag{10}$$

where $\mathbf{D}$ is a diagonal matrix, I is the identity matrix, and $tr(\mathbf{B})$ is the trace of matrix $\mathbf{B}$.

Proof. Kenneth Miller [14] proposed the following Lemma on Matrix Inverse.

$$(G+H)^{-1}=G^{-1}-\frac{1}{1+tr(GH^{-1})}G^{-1}HG^{-1} \tag{11}$$

Applying Miller's lemma, where $G=DAD$ and $H=\lambda I$, leads directly to Eq. 10. □

Mapping $\mathbf{D}$ to $\mathbf{W}^{i\frac{1}{2}}$ [1], $\mathbf{A}$ to either of $\mathbf{K}_X^i$ or $\mathbf{K}_Y^i$, we can compute $\mathcal{M}^i=\{{\mathbf{K}_X^i}^{-1},{\mathbf{K}_Y^i}^{-1}\}$. Having computed $\mathbf{W}^i$ on test time, $(\mathbf{W}^{i\frac{1}{2}}\mathbf{K}_X^i\mathbf{W}^{i\frac{1}{2}}+\lambda_x\mathbf{I})^{-1}$, $(\mathbf{W}^{i\frac{1}{2}}\mathbf{K}_X\mathbf{W}^{i\frac{1}{2}}+\lambda_x\mathbf{I})^{-1}$ could be computed in quadratic time given $\mathcal{M}^i$ following equation 10, since the inverse and the power of $\mathbf{W}^{i\frac{1}{2}}$ has linear computational complexity since it is diagonal.

5.2 Prediction

ODC-Prediction is performed in three steps.

(1) Finding the closest subdomains. The closest $K'\ll K$ subdomains are determined based on the covariance norm of the displacement of the test input from the means of the subdomain distribution (i.e. $\|\mathbf{x}-\boldsymbol{\mu}_i'\|_{\Sigma_i'^{-1}}, i=1:K$, where $\|\mathbf{x}-\boldsymbol{\mu}_i'\|_{\Sigma_i'^{-1}}=(\mathbf{x}-\boldsymbol{\mu}_i')^T\Sigma_i'^{-1}(\mathbf{x}-\boldsymbol{\mu}_i')$. The reason behind using the covariance norm is that it captures details of the density of the distribution in all dimensions. Hence, it better models $p(\mathbf{x}|D_i)$, indicating better prediction of $\mathbf{x}$ on D_i.

(2) Closest subdomains Prediction. Having determined the closest subdomains, predictions are made for each of the closest clusters. We denote these predictions as $\{Y_{x_*}^i\}_{i=1}^{K'}$. Each of these prediction are computed according to the selected kernel machine. For GPR, predictive mean and variance are $O(M\cdot d_X)$ and $O(M^2\cdot d_Y)$ respectively, for each output dimension. For TGP, the prediction is $O(l_2\cdot M^2\cdot d_Y)$; see Eq 2.

[1] $\mathbf{W}$ is a diagonal matrix

(3) Subdomains weighting and Final prediction. The final predictions are formulated as $\mathbf{Y}(x_*)=\sum_{i=1}^{K'}a_i\mathbf{Y}_{x_*}^i, a_i>0, \sum_{i=1}^{K'}a_i=1$. $\{a_i\}_{i=1}^{K'}$ are computed as follows. Let the distribution of domain $\{D_{x_*}^i=\|x-\mu_i'\|_{\Sigma_k'^{-1}}\}_{i=1}^{K'}$ denotes to the distances to the closest subdomains, $\{L_{x_*}^i=1/D_{x_*}^i\}_{i=1}^{K'}$, $a_i=L_{x_*}^i/\sum_{i=1}^{K'}L_{x_*}^i$.

It is not hard to see that when $K'=1$, the prediction step reduces to regression using the closest subdomain to the test point. However it is reasonable in most of the prior work to make prediction using the closest model, we generalized it to K' closest kernel machines and combining their predictions, so as to study how consistency of the combined prediction behaves as the overlap increases (i.e., p); see the experiments.

6 Experimental Results

Equal-Size Kmeans Step Experiment: We also tried another variant for Ekmeans that we call Iterative Minimum-Distance Assignments EKmeans (IMDA- Ekmeans). Note that the algorithm presented earlier in the paper is denoted as Assign and Balance Kmeans (AB-Kmeans). The IMDA-Ekmeans algorithm works as follows. We initialize a pool of unassigned points $\tilde{X}=X$ and initialize all clusters as empty. Given the means computed from the previous update steps, we compute the distances $d(\mathbf{x}_i,\mu_j)$ for all points/center pairs. We iteratively pick the minimum distance pair

$$(\mathbf{x}_p,\mu_l):d(\mathbf{x}_p,\mu_l)\leq d(\mathbf{x}_i,\mu_j)\forall x_i\in\tilde{X}\text{and}|C_l|<N/K$$

and assign point x_p to cluster l. The point is then removed from the pool of unassigned points. if $|C_l|=N/K$, then it is marked as balanced and no longer considered. The process is repeated until the pool is empty; see Algorithm 2.

Table 3 presents the average cost over 10 runs of IMDA-Ekmeans and AB-Ekmeans algorithms. We initialize both the AB-Ekmeans and IMDA-EKmeans algorithms by the cluster centers computed by running the standard k-means.

Input: $\mathbf{X}(N\times d_x), \{\boldsymbol{\mu}_i\}_{i=1}^K$
Output: labels
1- Create a matrix $\mathbf{D}\in R^{N\times K}$, where $\mathbf{D}[i,j]$ is the distance between the i^{th} point to the j^{th} cluster center.
2- Get the coordinate (i_*,j_*) that maps the smallest distance in $\mathbf{D}$.
3- Remove the i_*^{th} row from matrix $\mathbf{D}$ and mark it as assigned to the j^{th} cluster
4- If the size of the cluster j achieves the ideal size (i.e. n/K), then remove the j^{th} column from matrix $\mathbf{D}$.
5- Go to step 2 if there is still unassigned points

Algorithm 2: Iterative Minimum-Distance Assignments (IMDA) k-means: Assignment Step

As illustrated in table 3, the AB-Ekmeans outperforms IMDA-Ekmeans in these experiments, which motivated us to utilize AB Ekmeans, which is presented in the paper, against IMDA-Ekmeans under our ODC prediction framework. Our interpretation for these results is because AB-Ekmeans initializes the assignment with an assignment that minimizes the cost $J(C) = \min \sum_{j=1}^{K} \sum_{\mathbf{x_i} \in C_j} d(\mathbf{x_i}, \mu_j)$ given the cluster centers and then balance the clusters. In all the following experiments, we uses AB-EKmeans due to its clear superior performance to IMDA-EKmeans.

Table 3: $J(C)$ of AB-kmeans and IMDA-kmeans on a dataset of 10,000 random 2D points, averaged over 10 runs

	K=5	K = 10	K=50
AB-kmeans	1077.3	540.241	105.505
IMDA-kmeans	1290.6	657.446	122.006
Error Reduction	16.53%	17.83%	13.52%

	Poser Dataset (Agarwal etal, 06)	Human Eva Dataset (Sigal etal, 2006)	Human 3.6M Dataset (Lonescu etal, 2014)
Training Size	1,927	4,815	335 K
Test Size	418	4,815	165 K
Features	boW silhouette features	HOG	HOG
Error	$\frac{1}{54}\sum_{m=1}^{54}\|\hat{y}^m - y^{*m} mod\ 360°\|$	$\frac{1}{L}\sum_{m=1}^{L}\|\hat{y}^m - y^{*m}\|$	
Number of Views	1	3	4
Pose representation	Joint angles	Joints relative position	

$\hat{y}$ is the estimated pose, y^* is the ground truth pose

Fig. 5: Datasets, Representations, and Features

Datasets and Setup. We evaluated our framework on three human pose estimation datasets, Poser, HumanEva, and Human3.6M; see Fig. 5 for summary of setup and representation for each. Poser dataset [1] consists of 1927 training and 418 test images. The image features, corresponding to bag-of-words representation with silhouette-based shape-context features. The error is measured by the root mean-square error (in degrees), averaged over all joints angles, and is given by: $Error(\hat{y}, y^*) = \frac{1}{54}\sum_{m=1}^{54} \|\hat{y}^m - y^{*m} mod\ 360°\|$, where $\hat{y} \in R^{54}$ is an estimated pose vector, and $y^* \in R^{54}$ is a true pose vector. HumanEva datset [20] contains synchronized multi-view video and Mocap data of 3 subjects performing multiple activities. We use HOG features [7] ($\in R^{270}$) proposed in [2]. We use training and validations subsets of HumanEva-I and only utilize data from 3 color cameras with a total of 9630 image-pose frames for each camera. This is consistent with experiments in [2] and [26]. We use half of the data for training and half for testing. Human3.6M [4] is a dataset of millions of Human poses. We managed to evaluate our proposed ODC-framework on six Subjects (S1, S2, S6, S7, S8, S9) from it, which is ≈ 0.5 million poses. We split them into 67% training 33% is testing. HOG features are extracted for 4 image-views for each pose and concatenated into 3060-dim vector. Error for each pose, in both HEva (in mm) and Human 3.6 (in cm), is measured as $Error(\hat{y}, y^*) = \frac{1}{L}\sum_{m=1}^{L} \|\hat{y}^m - y^{*m}\|$.

There are four control parameters in our ODC framework: M, p, t, and K'. Figure 6 shows our parameter analysis with different values of p, t and K' on HumanEva dataset for GPR and TGP as local regression machines, where $M = 800$. Each sub-figure consists of six plots in two rows. The first row indicates the results using AB-Ekmeans clustering scheme, while the second row shows the results for RPC clustering scheme. Each row has three plots, one for $K' = 1$, 2, and 3 respectively. Each plot shows the error of different t against p from 0 to 0.95; i.e., it shows how the overlap affects the performance for different values of t. Each plot shows, on its top caption, the minimum and the maximum overlap regression errors where $t \to 1$. Looking at these plots, there are a number of observations:

(1) As $t \to 1$ (the solid red line), the behavior of the error tends to reduce as p increases, i.e., the overlap.

(2) Comparing different K', the behavior of the error indicates that combining multiple predictions (i.e., $K' = 2$ and $K' = 3$), gives poor performance, compared with $K' = 1$, when the overlap is small. However, all of them, $K' = 1$, 2, and 3, performs well as $p \to 1$; see column 2 and 3 in Fig. 6 and Fig. 8. This indicates consistent prediction of neighboring subdomains as p increases; see also Fig. 7 for side by side comparison of different K'. The main reason behind this behavior is that as p increases, the local models of the neighboring subdomains normally share more training points on their boundaries, which is reflected as shared constraints during the training of these models making them more consistent on prediction.

(3) Comparing the first row to the second row in each sub-figure, it is not hard to see that our AB-Ekmeans partitioning scheme consistently outperforms RPC [5], *e.g.* the error in cases of GPR (M=800) is 47.48mm for AB-EKmeans and 50.66mm for RPC, TGP (M=800) is 38.8mm for AB-EKmeans and 39.8mm for RPC. This problem is even more severe when using smaller M, *e.g.* the error in case of TGP (M=400) is 39.5mm for EKmeans and 47.5mm for RPC; see a detailed plot for M=400 in Fig. 9. We noticed sigficant drop in the performance as M decreases. For instance when $M = 200$, The error for TGP best performance increased to 43.88mm instead of 38mm for $M = 800$.

(4) TGP gives better prediction than GPR (i.e., 38mm using TGP compared with 47mm using GPR).

(5) As M increases, the prediction error decreases. For instance, when $M = 200$, The error for TGP best performance increased to 43.88mm instead of 38.9mm for $M = 800$.

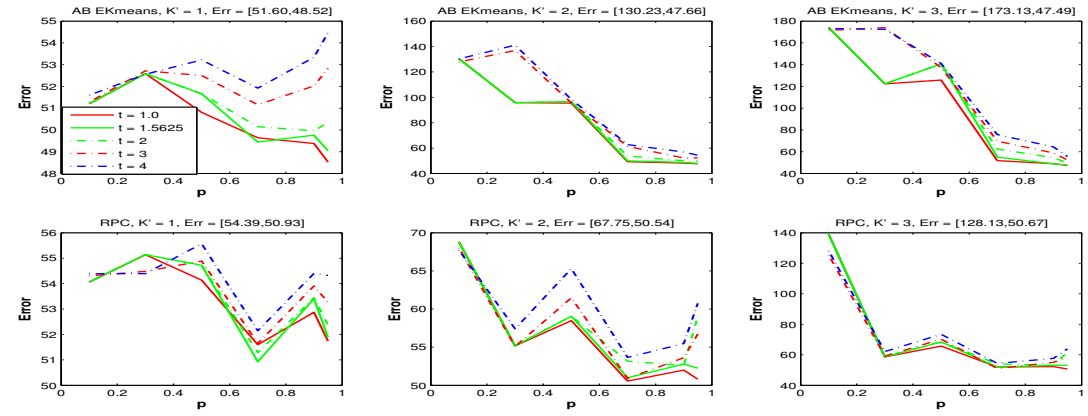
AB EKmeans, K' = 1, Err = [51.60,48.52]
AB EKmeans, K' = 2, Err = [130.23,47.66]
AB EKmeans, K' = 3, Err = [173.13,47.49]
RPC, K' = 1, Err = [54.39,50.93]
RPC, K' = 2, Err = [67.75,50.54]
RPC, K' = 3, Err = [128.13,50.67]
t = 1.0
t = 1.5625
t = 2
t = 3
t = 4
Error
p

(a) GPR-ODC (M=800)

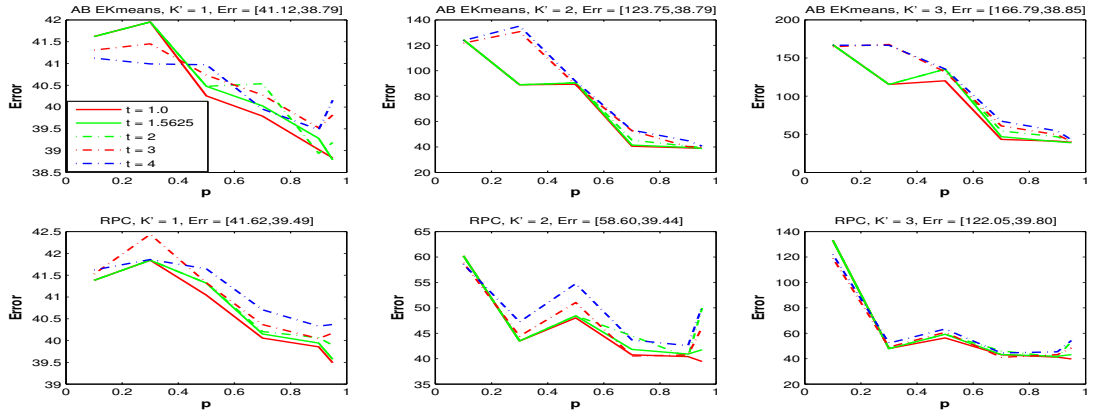
AB EKmeans, K' = 1, Err = [41.12,38.79]
AB EKmeans, K' = 2, Err = [123.75,38.79]
AB EKmeans, K' = 3, Err = [166.79,38.85]
RPC, K' = 1, Err = [41.62,39.49]
RPC, K' = 2, Err = [58.60,39.44]
RPC, K' = 3, Err = [122.05,39.80]
t = 1.0
t = 1.5625
t = 2
t = 3
t = 4
Error
p

(b) TGP-ODC (M=800)

Fig. 6: ODC framework Parameter Analysis of GPR and TGP on Human Eva Dataset

Table 4: Error & Time on Poser and Human Eva datasets (Intel core-i7 2.6GHZ), M = 800

		Poser			**HumanEva**		
		Error (deg)	**Training Time**	**Prediction Time**	**Error (mm)**	**Training Time**	**Prediction Time**
TGP	**NN [2]**	5.43	-	188.99 sec	**38.1**	-	6364 sec
	ODC ($p = 0.9, t = 1, K' = 1$)**-Ekmeans**	**5.40**	(3.7 +25.1) sec	**16.5** sec	**38.9**	(2001 + 45.4) sec	**298** sec
	ODC ($p = 0, t = 1, K' = 1$)**-Ekmeans**	7.60	(3.9 + 1.33) sec	14.8 sec	41.87	(240 + 4.9) sec	257 sec
	ODC ($p = 0.9, t = 1, K' = 1$)**-RPC**	5.60	(0.23 +41.6) sec	15.8 sec	39.9	(0.45 + 49.1) sec	277 sec
	ODC ($p = 0, t = 1, K' = 1$)**-RPC**	7.70	(0.15 + 1.7) sec	13.89 sec	42.32	(0.19 + 5.2) sec	242 sec
GPR	**NN**	6.77	-	24 sec	54.8	-	618 sec
	ODC ($p = 0.9, t = 1, K' = 1$)**-Ekmeans**	**6.27**	(3.7 +11.1) sec	**0.56** sec	**49.3**	(2001 + 42.85)sec	**79** sec
	ODC($p = 0.0, t = 1, K' = 1$)**-Ekmeans**	7.54	(3.9 + 1.38 sec)	0.35 sec	49.6	(240 + 6.4) sec	48 sec
	ODC ($p = 0.9, t = 1, K' = 1$)**-RPC**	6.45	(0.23 +17.3) sec	0.52 sec	52.8	(0.49 + 46.06) sec	64 sec
	ODC ($p = 0.0, t = 1, K' = 1$)**-RPC = [5]**	7.46	(0.15 + 1.5) sec	0.27 sec	54.6	(0.26 + 4.6) sec	44 sec
	FIC [22]	7.63	(- + 20.63)	0.3106	68.36	-	102 sec

We found these observation to be also consistent on Poser dataset.

This analysis helped us conclude recommending choosing t close to 1, big overlap (p closer to 1), and $K' = 1$ is sufficient for accurate prediction.

Having accomplished the performance analysis which comprehensively interprets our parameters, we used the recommended setting to compare the performance with other methods and show the benefits of this framework. Figure 10 shows the speedup gained by retrieving the matrix inverses on test time, compared with computing them at test time by NN scheme. The figure shows significant speedup from precomputing local kernel machines.

Table 4 shows error, training time and prediction time of NN, FIC, and different variations of ODC on Poser and Human-Eva datasets. Training time is formatted as (t_c + t_p), where t_c is the clustering time and t_p is the remaining training time excluding clustering. As indicated in the top part of table 4, TGP under our ODC-framwork can significantly speedup the prediction compared with NN-scheme

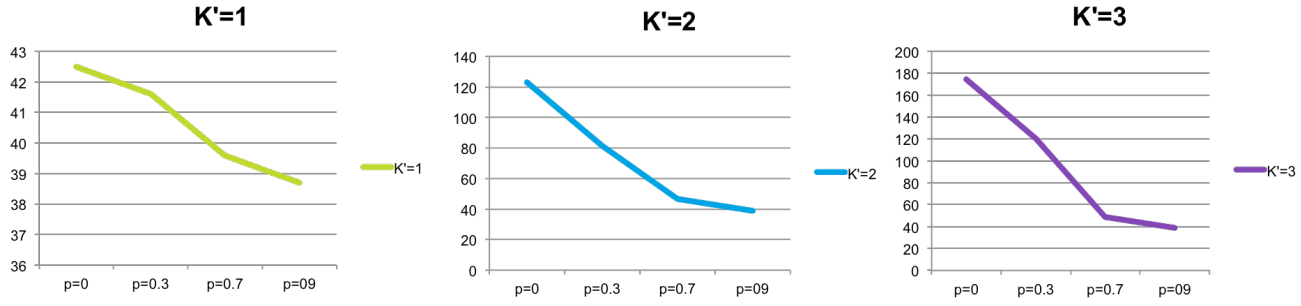

Fig. 7: HumanEva TGP different K' as overlap increase, , M=800

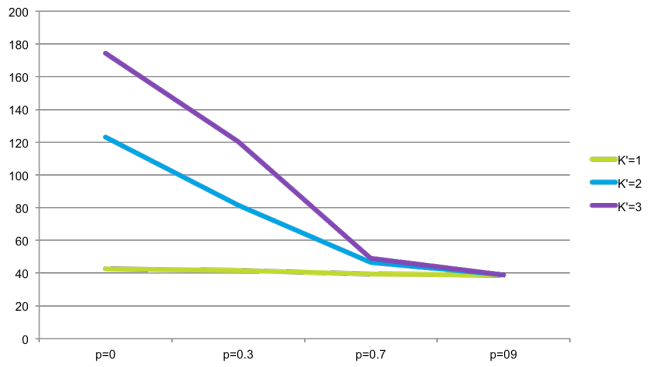

Fig. 8: Increasing K' significantly heart the performance for small overlap (Human Eva TGP, M=800)

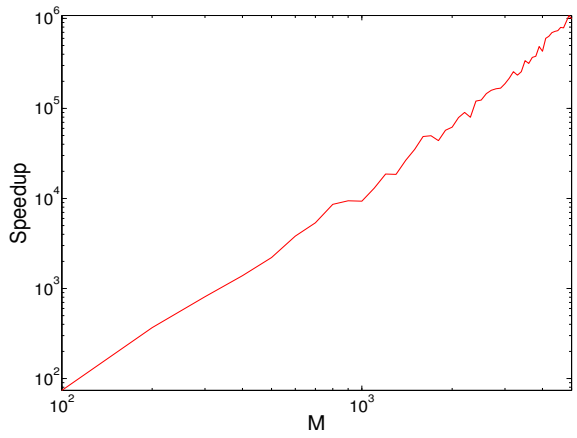

Fig. 10: Speedup of ODC framework prediction on either TGP or GPR while retrieving precomputed matrix inverses as M increases, compared with computing them on test time by KNN scheme (log-log scale)

in [2], while achieving competitive performance; better in case Poser Dataset. As illustrated in our analysis in Figure 6, higher overlap (p) gives better performance. From time analysis perspective, higher p costs more training time due that more subdomains are created and trained. While, Figure 6 and Table 4 indicates that AB-Ekmeans gives better performance than RPC under both GPR and TGP, AB-Ekmeans takes more time for clustering. Yet, it is feasible to compute in all the datasets, we used in our experiments. Our experiments also indicate that as $p \to 1$ in TGP and GPR, $K' = 2$ and $K' = 3$ takes double and triple the prediction time respectively, compared with $K' = 1$, with almost no error reduction. We also compared our model to FIC in case of GPR, and our model achieved smaller error and smaller prediction time; see bottom part in Table 4. However, TGP consistently gives better results on both Poser and HumanEva datasets. We also tried full TGP and GPR on Poser and Human Eva Datasets. Full TGP error is 5.35 for Poser and 40.3 for Human Eva. Full GPR error is 6.10 for Poser and 59.62 for Human Eva. The results indicate that ODC achieves either better or competitive to the full models. Meanwhie, the speedup is sigbificant for TGP prediction ($21X$ for Human Eva and $11X$ for Poser Datasets); see Fig. 11. For GPR prediction, we achieved the best performance and the lowest prediction time compared to existing GPR prediction methods; see Fig. 12.

Based on our comprehensive experiments on HumanEva and Poser datasets, we conducted an experiment on Human3.6M dataset with TGP kernel machine, where $M = 1390$, $t = 1$, $p = 0.6, K' = 1$, Ekmeans for clustering. We achieved a speedup of $41.7X$ on prediction time using our ODC framework compared with NN-scheme, i.e., 7 days if NN-scheme is used versus 4.03 hours in our case with our MATLAB implementation. The error is 13.5 (cm) for NN and 13.8 (cm) for ODC; see Fig. 13.

7 Conclusion

We proposed an efficient ODC framework for kernel machines and validated the framework on structured regression machines on three human pose estimation datasets. The key idea is to equally partition the data and create cohesive overlapping subdomains, where local kernel machines are computed for each of them. The framework is general and could be applied to various kernel machine beyond GPR, TGP, IWTGP validated in this work. Similar to TGP and IWTGP,

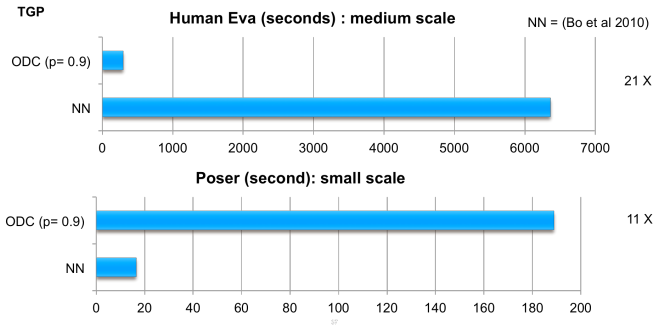

Fig. 11: TGP Human Eva Dataset (Speed)

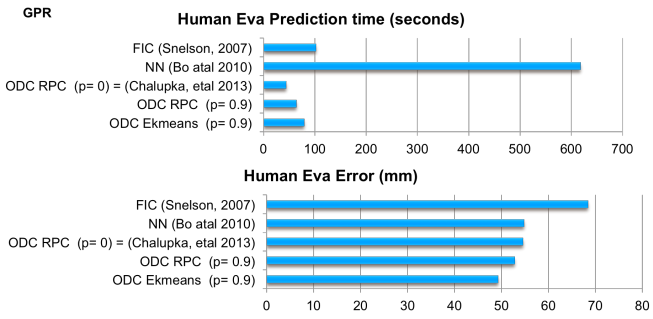

Fig. 12: GPR speed and error (Human Eva Dataset)

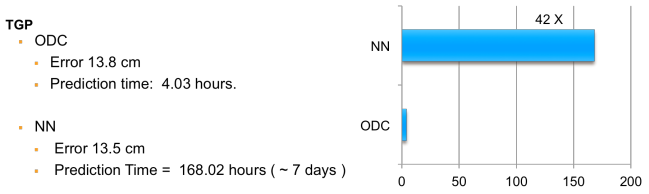

Fig. 13: TGP speed and error (Human3.6M Dataset)

our framework could be easily applied to the recently proposed Generalized TGP [8] which is based on Sharma Mittal divergence, a relative entropy measure brought from Physics community. We also theoretically justified our framework's notion.

Acknowledgment. This research was partially funded by NSF award # 1409683.

References

1. Agarwal, A., Triggs, B.: Recovering 3d human pose from monocular images. PAMI (2006)
2. Bo, L., Sminchisescu, C.: Twin gaussian processes for structured prediction. Int. J. Comput. Vision **87**(1-2), 28–52 (2010)
3. Bonchi, F., Gionis, A., Ukkonen, A.: Overlapping correlation clustering. Knowledge and information systems **35**(1), 1–32 (2013)
4. Catalin Ionescu Dragos Papava, V.O., Sminchisescu, C.: Human3.6M: Large Scale Datasets and Predictive Methods for 3D Human Sensing in Natural Environments. Tech. rep., nstitute of Mathematics at the Romanian Academy and University of Bonn (2012)
5. Chalupka, K., Williams, C.K.I., Murray, I.: A framework for evaluating approximation methods for gaussian process regression. JMLR **14**(1) (2013)
6. Cristianini N. Shawe-Taylor, J., Kandola, J.S.: Spectral kernel methods for clustering. In: NIPS (2001)
7. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: CVPR, vol. 1, pp. 886 –893 vol. 1 (2005). DOI 10.1109/CVPR.2005.177
8. Elhoseiny, M., Elgammal, A.: Generalized twin gaussian processes using sharma-mittal divergence. Machine Learning pp. 1–26 (2015)
9. Gramacy, R.B., Lee, H.K.H.: Bayesian treed gaussian process models with an application to computer modeling. Journal of the American Statistical Association (2007)
10. Gretton, A., Bousquet, O., Smola, A., Schölkopf, B.: Measuring statistical dependence with hilbert-schmidt norms. In: International conference on Algorithmic Learning Theory (2005)
11. Hartigan, J.A., Wong, M.A.: Algorithm as 136: A k-means clustering algorithm. J. Roy. Statist. Soc. Ser. C) (1979)
12. Hoerl, A.E., Kennard, R.W.: Ridge Regression: Biased Estimation for Nonorthogonal Problems. Technometrics (1970)
13. Khandekar, R., Kortsarz, G., Mirrokni, V.: On the advantage of overlapping clusters for minimizing conductance. Algorithmica **69**(4), 844–863 (2014)
14. Miller, K.S.: On the inverse of the sum of matrices. Mathematics Magazine **54**(2), 67–72 (1981)
15. Muja, M., Lowe, D.: Fast approximate nearest neighbors with automatic algorithm configuration. In: VISAPP, pp. 331–340 (2009)
16. Park, C., Huang, J.Z., Ding, Y.: Domain decomposition approach for fast gaussian process regression of large spatial data sets. Journal of Machine Learning Research **12**, 1697–1728 (2011)
17. Pérez-Suárez, A., Martínez-Trinidad, J.F., Carrasco-Ochoa, J.A., Medina-Pagola, J.E.: Oclustr: A new graph-based algorithm for overlapping clustering. Neurocomputing **121**, 234–247 (2013)
18. Rasmussen, C.E., Williams, C.K.I.: Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning). The MIT Press (2005)
19. Seeger, M.: Fast forward selection to speed up sparse gaussian process regression. In: in Workshop on AI and Statistics 9 (2003)
20. Sigal, L., Balan, A.O., Black, M.J.: Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion
21. Snelson, E.: Local and global sparse gaussian process approximations. In: AISTATS (2007)
22. Snelson, E., Ghahramani, Z.: Sparse gaussian processes using pseudo-inputs. In: NIPS (2006)
23. Tresp, V.: A bayesian committee machine. NEURAL COMPUTATION **12**, 2000 (2000)
24. Urtasun, R., Darrell, T.: Sparse probabilistic regression for activity-independent human pose inference. In: CVPR (2008)
25. Williams, C., Seeger, M.: Using the nystrm method to speed up kernel machines. In: NIPS (2001)
26. Yamada, M., Sigal, L., Raptis, M.: No bias left behind: covariate shift adaptation for discriminative 3d pose estimation. In: ECCV (2012)
27. Yamada, M., Suzuki, T., Kanamori, T., Hachiya, H., Sugiyama, M.: Relative density-ratio estimation for robust distribution comparison. In: NIPS, pp. 594–602 (2011)
28. Yuksel, S., Wilson, J., Gader, P.: Twenty years of mixture of experts. Neural Networks and Learning Systems, IEEE Transactions on (2012)
29. Zhu, S., Wang, D., Li, T.: Data clustering with size constraints. Knowledge-Based Systems **23**(8), 883 – 889 (2010)

Appendices

A IWTGP-ODC Experiments

Tables 5 and 6 details the results of IWTGP-ODC experiments on Poser and HumanEva datasets in terms of error and speedup in prediction time.

	IWTTGP ($M = 800, M_{tst} = 418$)	**IWTGP-ODC** ($M = 800, M_{tst} = 418$)
error (deg)	6.1	5.32
err reduction (deg)	-	**0.783**
err reduction %	-	**12.836%**
Prediction Time (sec)	360.0	26.61
speedUp	-	**13.5**

Table 5: POSER dataset IWTGP-NN vs IWTGP-ODC

	IWTGP ($M = M_{tst} = 800$)	**IWTTGP-ODC** ($M = M_{tst} = 800$)
error (mm)	39.1	39.3
err reduction (mm)	-	**-0.2**
err reduction %	-	**-0.512%**
Prediction Time (sec)	7938.15	569.66
speedUp	-	**13.92**

Table 6: Humen Eva dataset: IWTGPKNN vs IWTGP-ODC

B More figures on AB Ekmeans

Figure 14 shows the clustering performance on 300000 random 2D point (K=5). Figure 15 shows the clustering output of our algorithm visualized on using the first three principal components of Human Eva training hog features. The figures shows that the cluster are spatially cohesive but not necessarily circular. This makes the elliptic distribution of the data captured by Mode 3 gives more accuracy membership measure me to the subdomains.

C Overlapping Domain Cover(ODC) Generation-Algorithm

Algorithm 3 shows how the overlapping sub-domains are generated form the the equal size clusters from the closest r clusters.

D Local Kernel Machines hyper-parameters on each dataset

The hyper parameters were learnt using cross validation on the training set for GPR, TGP and IWTGP that we are interested in. The following subsection present the learnt hyper-parameters and the error measures on each dataset in case of TGPs.

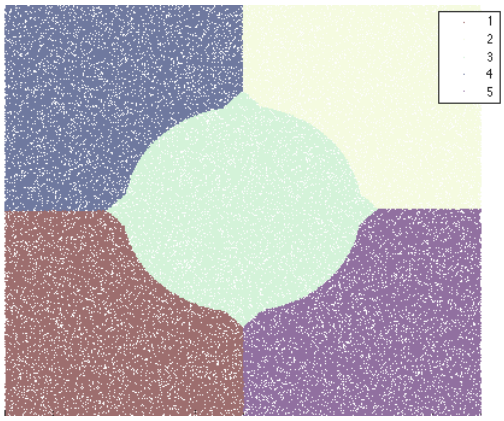

(a) 5 clusters

Fig. 14: Applying our Assign and Balance variant of Kmeans on 300,000 random 2D points

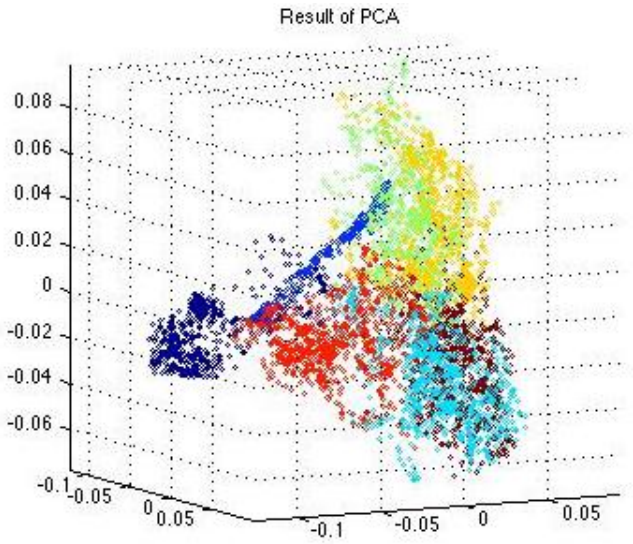

Fig. 15: Human Eva clustering first three Pricipal Components

D.1 Poser Dataset

The parameters $2\rho_x^2$, $2\rho_y^2$, λ_X, and λ_Y were assigned to 5, 5000, 10^{-4}, and 10^{-4}, respectively.

D.2 HumanEva Dataset

The parameters $2\rho_x^2$, $2\rho_y^2$, λ_X, and λ_Y were assigned to 5, 500000, 10^{-3}, and 10^{-3}, respectively.

Input: Clusters $\{C_k\}_{k=1}^{K}$ **Output**: Overlapping subdomains $\{D_k\}_{k=1}^{K}$
foreach *Cluster* C_k **do**
Compute the closest r clusters $\{C^{'}{}_i\}_{i=1}^{r}$ based on $DK_i = \|\mu_k - \mu_i\|, i \neq k$
Let $LK_i = 1/DK_i, WK_i = \frac{LK_i}{\sum_{l=1}^{OCC} LK_l}\ i = 1:r$
Let $NPK_i = floor(WK_i * OPC), i = 1:r$
Let $ExKPts = (1-p)M - \sum_{l=1}^{r} NPK_l$
Let NPK_i = $NPK_i + 1$, $i = 1 : ExKPts$
$D_k = C_k$
Let $overflow = 0$
▷ The following for loop goes over the r clusters on an increasing order of DK_i
for *i=1 : r* **do**
if NPK_i¿ $|C_i|$ **then**
$overflow = overflow + NPK_i - |C_i|$
$NPK_i = |C_i|$
if NPK_i¡ $|C_i|$ **then**
$G_i = min(overflow, |C_i| - NPK_i)$
$NPK_i = NPK_i + G_i$
$overflow = overflow - G_i$
$Ps_i = KNN(OVC_{Kj}, NPK_i)$
$D_k = D_k \cup Ps_i$
for *i=1 : r* **do**
$Ps_i = KNN(OVC_{Kj}, NPK_i)$
$D_k = D_k \cup Ps_i$
▷ where KNN is the K-nearest neighbors algorithms. For high performance calculation of KNN, we use FLANN [15] to calculate KNN.

Algorithm 3: Subdomains Generation (Note: All $\{D_k\}_{k=1}^{K}$ are stored as indices to X).

D.3 Human 3.6 Dataset

The parameters $2\rho_x^2$, $2\rho_y^2$, λ_X, and λ_Y were assigned to 5, 500000, 10^{-3}, and 10^{-3}, respectively.

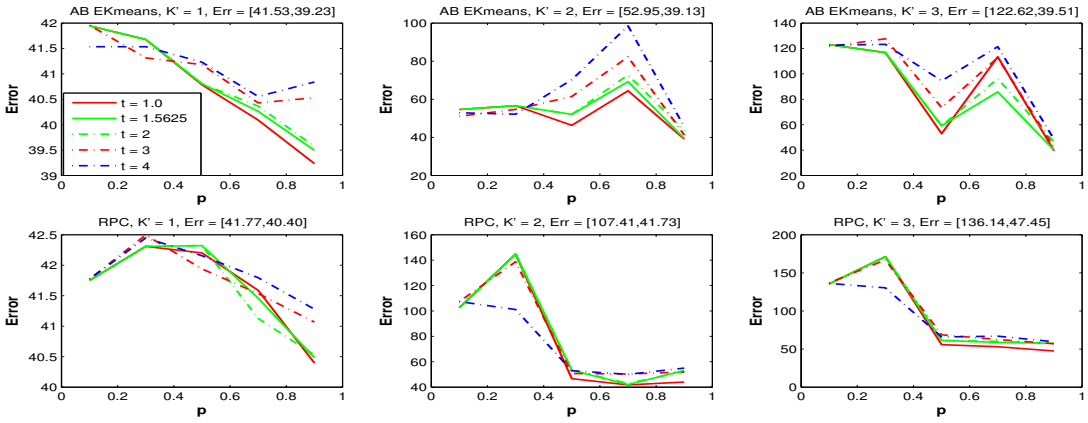
AB EKmeans, K' = 1, Err = [41.53,39.23]
AB EKmeans, K' = 2, Err = [52.95,39.13]
AB EKmeans, K' = 3, Err = [122.62,39.51]
RPC, K' = 1, Err = [41.77,40.40]
RPC, K' = 2, Err = [107.41,41.73]
RPC, K' = 3, Err = [136.14,47.45]
t = 1.0
t = 1.5625
t = 2
t = 3
t = 4
Error
p

(a) TGP-ODC (M=400)

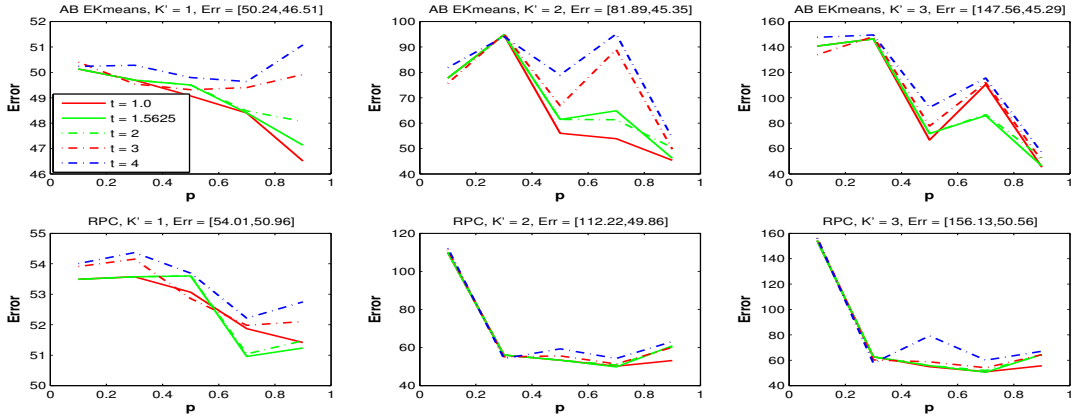
AB EKmeans, K' = 1, Err = [50.24,46.51]
AB EKmeans, K' = 2, Err = [81.89,45.35]
AB EKmeans, K' = 3, Err = [147.56,45.29]
RPC, K' = 1, Err = [54.01,50.96]
RPC, K' = 2, Err = [112.22,49.86]
RPC, K' = 3, Err = [156.13,50.56]
t = 1.0
t = 1.5625
t = 2
t = 3
t = 4
Error
p

(b) GPR-ODC (M=400)

Fig. 9: Overlapping Domain Cover Parameter Analysis of GPR and TGP on Human Eva Dataset (best seen in color) (M=400)

Mohamed Elhoseiny is a PostDoc Researcher at Facebook Research. His primary research interest is in computer vision, machine learning, intersection between natural language and vision, language guided visual-perception, and visual reasoning, art & AI. He received his PhD degree from Rutgers University, New Brunswick, in 2016 under Prof. Ahmed Elgammal. Mohamed received an NSF Fellowship in 2014 for the Write-a-Classifier project (ICCV13), best intern award at SRI International 2014, and the Doctoral Consortium award at CVPR 2016.

Ahmed Elgammal is a professor at the Department of Computer Science, Rutgers, the State University of New Jersey Since Fall 2002. Dr. Elgammal is also a member of the Center for Computational Biomedicine Imaging and Modeling (CBIM). His primary research interest is computer vision and machine learning. His research focus includes human activity recognition, human motion analysis, tracking, human identification, and statistical methods for computer vision. Dr. Elgammal received the National Science Foundation CAREER Award in 2006. Dr. Elgammal has been the Principal Investigator and Co-Principal Investigator of several research projects in the areas of Human Motion Analysis, Gait Analysis, Tracking, Facial Expression Analysis and Scene Modeling; funded by NSF and ONR. Dr. Elgammal is Member of the review committee/board in several of the top conferences and journals in the computer vision field. Dr. Elgammal received his Ph.D. in 2002 from the University of Maryland, College Park. He is a senior IEEE member.